LSF Analyzer User's Guide

Second Edition, August 1998

Platform Computing Corporation

LSF Analyzer User's Guide

Revision Information for LSF Analyzer User's Guide

Edition	Description
First	This document describes LSF Analyzer 3.1
Second	Revised to reflect the changes in LSF Analyzer 3.2

Contents

Preface

Audience

This guide is designed to help managers, LSF cluster administrators and performance analysts use LSF Analyzer to perform accounting and chargeback functions, cluster performance analysis, capacity planning and forecasting.

This guide describes the use of LSF Analyzer. In it, you will find all the information you need to set up and use LSF Analyzer at your site. This guide assumes you have knowledge of common LSF system terminology.

LSF Suite 3.2

LSF is a suite of workload management products including the following:

LSF Batch is a batch job processing system for distributed and heterogeneous environments, which ensures optimal resource sharing.

LSF JobScheduler is a distributed production job scheduler that integrates heterogeneous servers into a virtual mainframe or virtual supercomputer

LSF MultiCluster supports resource sharing among multiple clusters of computers using LSF products, while maintaining resource ownership and cluster autonomy.

LSF Analyzer is a graphical tool for comprehensive workload data analysis. It processes cluster-wide job logs from LSF Batch and LSF JobScheduler to produce

statistical reports on the usage of system resources by users on different hosts through various queues.

LSF Parallel is a software product that manages parallel job execution in a production networked environment.

LSF Make is a distributed and parallel Make based on GNU Make that simultaneously dispatches tasks to multiple hosts.

LSF Base is the software upon which all the other LSF products are based. It includes the network servers (LIM and RES), the LSF API, and load sharing tools.

There are two editions of the LSF Suite:

LSF Enterprise Edition

Platform's LSF Enterprise Edition provides a reliable, scalable means for organizations to schedule, analyze, and monitor their distributed workloads across heterogeneous UNIX and Windows NT computing environments. LSF Enterprise Edition includes all the features in LSF Standard Edition (LSF Base and LSF Batch), plus the benefits of LSF Analyzer and LSF MultiCluster.

LSF Standard Edition

The foundation for all LSF products, Platform's Standard Edition consists of two products, LSF Base and LSF Batch. LSF Standard Edition offers users robust load sharing and sophisticated batch scheduling across distributed UNIX and Windows NT computing environments.

Related Documents

The following guides are available from Platform Computing Corporation:

> *LSF Installation Guide*
> *LSF Batch Administrator's Guide*
> *LSF Batch Administrator's Quick Reference*
> *LSF Batch User's Guide*
> *LSF Batch User's Quick Reference*
> *LSF JobScheduler Administrator's Guide*
> *LSF JobScheduler User's Guide*
> *LSF Analyzer User's Guide*
> *LSF Parallel User's Guide*
> *LSF Programmer's Guide*

Online Documentation

- Man pages (accessed with the `man` command) for all commands
- Online help available through the Help menu for the `xlsbatch`, `xbmod`, `xbsub`, `xbalarms`, `xbcal`, `xlsjs`, `xlsadmin`, and `xanalyzer` applications.

Technical Assistance

If you need any technical assistance with LSF, please contact your reseller or Platform Computing's Technical Support Department at the following address:

> LSF Technical Support
> Platform Computing Corporation
> 3760 14th Avenue
> Markham, Ontario
> Canada L3R 3T7

> Tel: +1 905 948 8448
> Toll-free: 1-87PLATFORM (1-877-528-3676)
> Fax: +1 905 948 9975
> Electronic mail: *support@platform.com*

Please include the full name of your company.

You may find the answers you need from Platform Computing Corporation's home page on the World Wide Web. Point your browser to *www.platform.com*.

Preface

If you have any comments about this document, please send them to the attention of LSF Documentation at the address above, or send email to *doc@platform.com*.

1. Introduction

The LSF system provides a powerful distributed computing environment that tightly integrates a suite of products: LSF Batch, LSF JobScheduler, LSF MultiCluster, LSF Analyzer, LSF Make, LSF Parallel, and LSF Base. While each of these products independently delivers great value collectively they constitute a complete workload management solution. LSF Analyzer is a tool for comprehensive workload and performance analysis.

Overview of LSF Analyzer

LSF Analyzer processes historical workload data to produce reports about a cluster. The workload data includes information about batch jobs, system metrics, load indices and resource usage. LSF Analyzer provides system administrators and managers with the information to make intelligent, informed scheduling and capacity planning decisions required to fully utilize the power delivered by LSF.

LSF Analyzer can also be used to do chargeback accounting, generating chargeback reports and invoices.

The primary features of LSF Analyzer:

- Profiles highlighting the number of jobs processed by the system, job resource usage, system metrics, load indices and resource usage

- Usage trends for the LSF system hosts, users, queues, applications, and projects

- Information to manage resources by user and project

- Chargeback accounting for users or projects providing reports and invoices

- Data export to comma separated values (.csv) file format compatible with industry standard spreadsheet and data analysis tools

- Built-in and user-generated templates to automate analysis.

Basic Concepts

LSF Analyzer collects and analyzes historical data stored in the LSF database to produce statistical reports which are designed to suit your needs. The analysis can be displayed in table, bar, area and line charts, and can be saved as a template which makes it convenient to repeat the analysis any time.

The basic concepts used by LSF Analyzer.

- **Data Collection** The LSF data collection engine is fully integrated in the LSF system. During normal operation of the LSF system, historical data is collected for all LSF objects (jobs, users, queues, hosts, projects, load indices, and resources) over a user-determined period of time (hours, days, weeks, and months) and stored in the LSF database.

- **LSF Database** The LSF system works with commercial class database management systems (DBMS), providing superior performance and data management. The installation, configuration and maintenance of the LSF databases is discussed in *Chapter 6, 'LSF Database - UNIX', on page 41* and *Chapter 7, 'LSF Database - Windows NT', on page 49.*

- **LSF Analyzer** (xanalyzer) LSF Analyzer provides **xanalyzer**, a graphical analysis and reporting tool, as an integral part of this application. The xanalyzer application retrieves the stored data and performs statistical analysis to produce reports describing the LSF cluster and objects profiles.

Case Studies

The major advantage in using LSF Analyzer is it allows the LSF administrator to solve problems regarding the performance of the LSF cluster that would typically be very difficult to answer. Finding these solutions allows an LSF cluster to be configured and used optimally. Statistics generated by LSF Analyzer are used to show how well a system is working, and trend analysis helps with capacity planning.

Examples showing the benefits of LSF Analyzer:

- Who are the largest consumers of system resources?

- Are these resources being used efficiently?

- Are the service commitment levels being met (i.e., what is the clusters reliability)?

- What are the activity trends of a cluster?

User Profile

Who are the largest consumers of system resources?

LSF Analyzer can be used to identify the users who are submitting CPU-intensive jobs or are submitting a large number of jobs. With this type of information the administrator can take corrective action to prevent these users from monopolizing cluster resources.

1 Introduction

The report in *Figure 1* shows the number of jobs submitted by each user and the CPU resources consumed by these jobs.
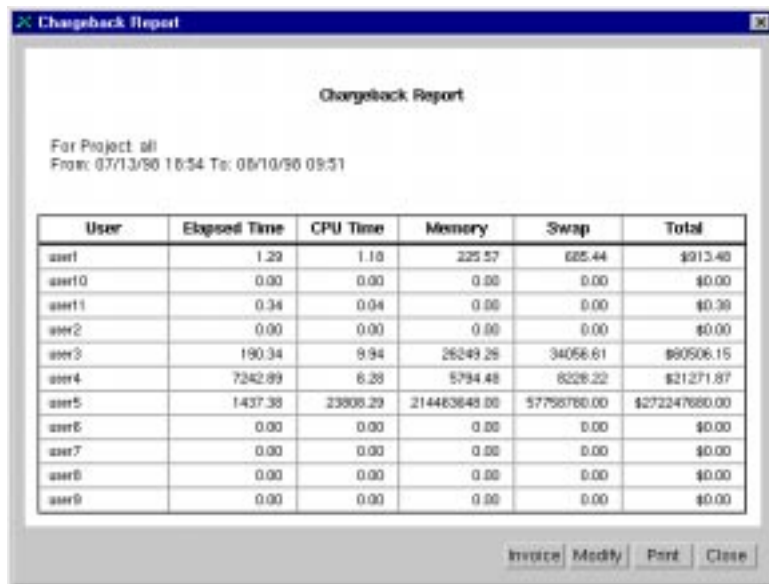
Figure 1. System Resource Usage
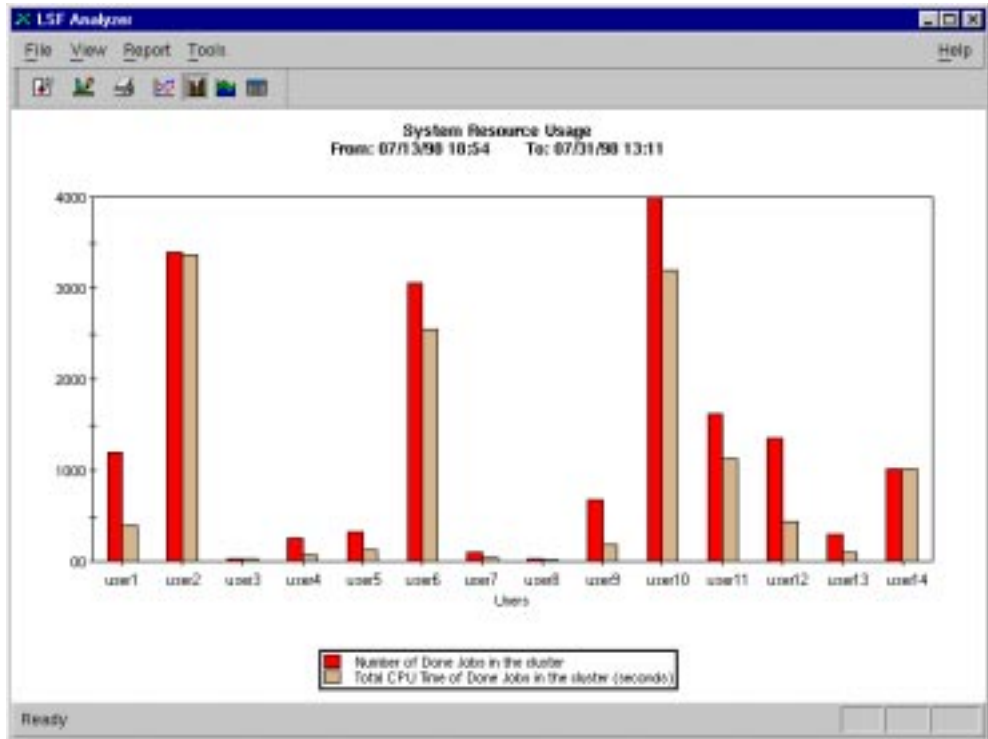
This report identifies the heaviest consumers of the LSF system resources. One possible action from this information is to implement fairshare policies to reduce the monopolization of system resources.

An extension of the user profile is the chargeback report shown in *Figure 2*.

Figure 2. Chargeback Report

Host Profile

Are the Host Resources being used efficiently?

LSF Analyzer can produce a host profile providing the justification to upgrade computing resource. Performance exceptions are identified in the cluster, like hosts that are not doing the expected amount of work due to hardware or configuration problems.

The report in *Figure 3* compares the memory utilization and CPU utilization for host1.

Figure 3. Host Resource Utilization
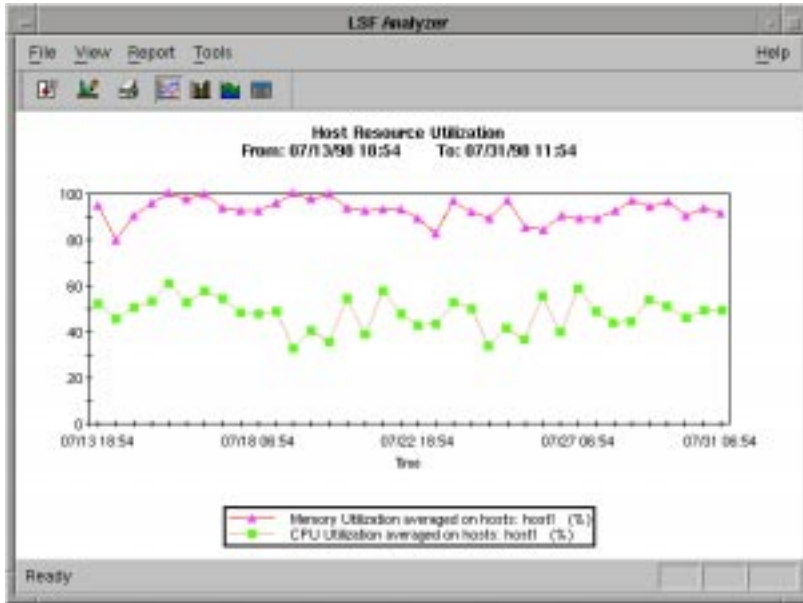
This example shows the memory for host1 is fully utilized but the available CPU resources are not. One conclusion drawn from this report is that there is not enough memory installed in host1. One possible corrective action would be to install additional memory resources, then rerun this report to verify CPU resources are being fully utilized.

Cluster Profile

Cluster Availability

Using LSF Analyzer to produce a cluster profile provides the information needed to demonstrate that service commitment levels were met.

The report shown in *Figure 4* was produced using the `/Performance/General/ClusterAvail_Time` template.

Figure 4. Cluster Availability
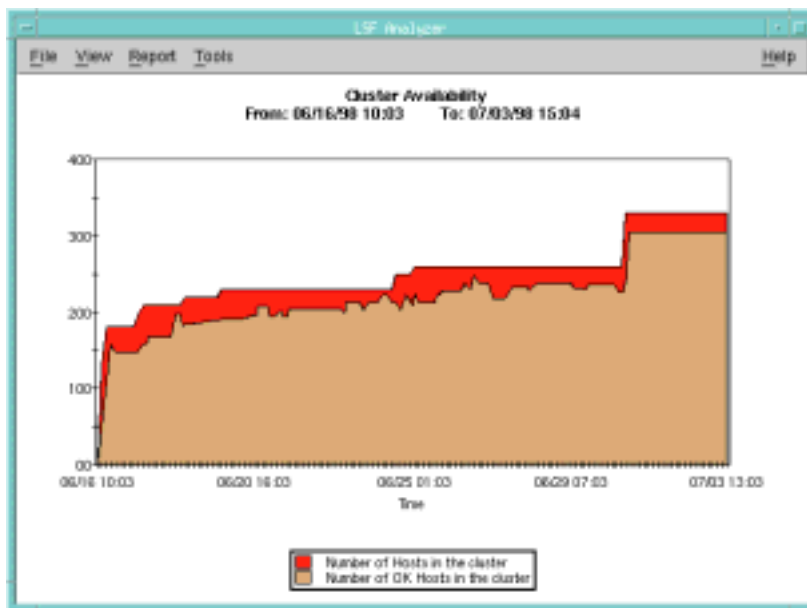
This report shows the number of OK hosts (available) in the cluster and the number of hosts (total) in the cluster. The large number of available (OK) hosts reflect the reliability of the cluster.

Activity Trend

Using LSF Analyzer to produce a cluster profile highlighting user submission trends provides the information needed to schedule regular maintenance and system downtime.

The report shown in *Figure 5* was produced using the `/Workload/General/Job_Time` template.

Figure 5. Number of Running and Suspended Jobs in the Cluster
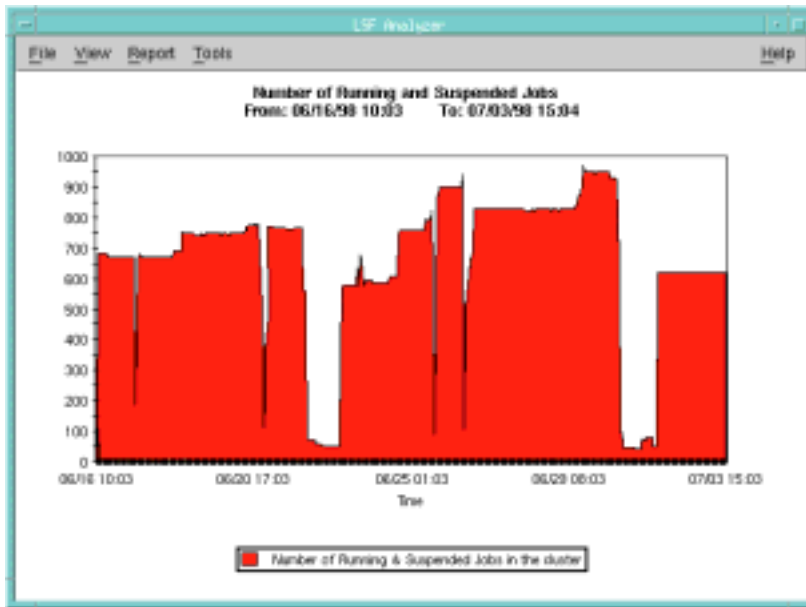
This report shows the number of running and suspended jobs in the cluster which represents the times of maximum and minimum system usage. This example shows that patterns of low system usage occur on a regular basis.

2. Getting Started

This chapter functions as a quick-start guide to begin using LSF Analyzer. The procedures in this chapter will provide you with experience using the various features of LSF Analyzer, by guiding you through the series of steps needed to generate a report.

A detailed explanation of the features and options for performing workload analysis is dealt with in the following chapters. Chargeback accounting is discussed in *Chapter 5 , Chargeback Accounting, on page 33.*

Using LSF Analyzer for the First Time

In this example, you will use the default template to create a report showing CPU usage for all users.

Configuring and Starting LSF Analyzer

Step 1   Start LSF Analyzer, displays **Welcome to LSF Analyzer** dialog.
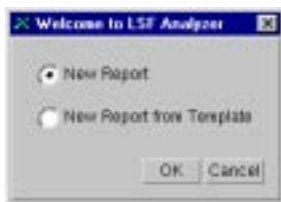
On UNIX enter the following command:

```
% xanalyzer
```

Note

> *If the command cannot be executed, the appropriate directories may need to be added to the systems path; check with the system administrator.*

On Windows NT select:

- **Programs | LSF Suite for Workload Management | LSF Analyzer**

Step 2   Select **New Report**, click **OK**, displays **Select Data Source** dialog.

Step 3   Select a data source, click **OK**, displays **New Report** dialog.
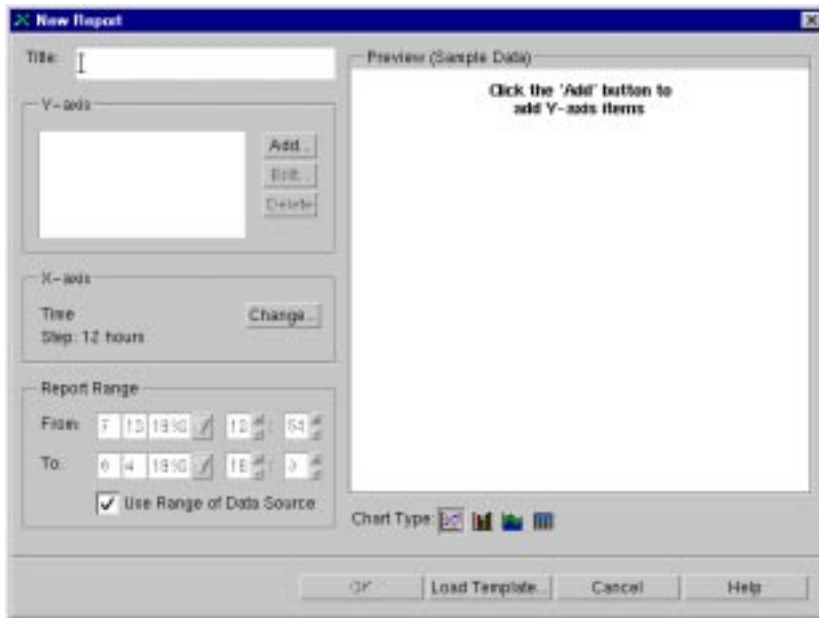
Step 4   Enter *Number of Jobs* as the **Title** for this Report.

Step 5 Click Y-axis **Add...** button, displays **Add to Y-axis** dialog.
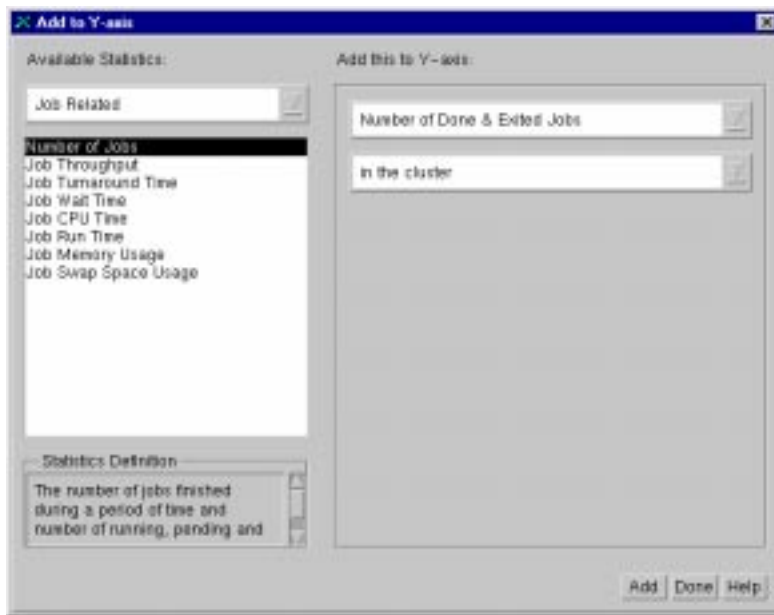
Step 6 Click **Add** then **Done**, returns to **New Report** dialog.
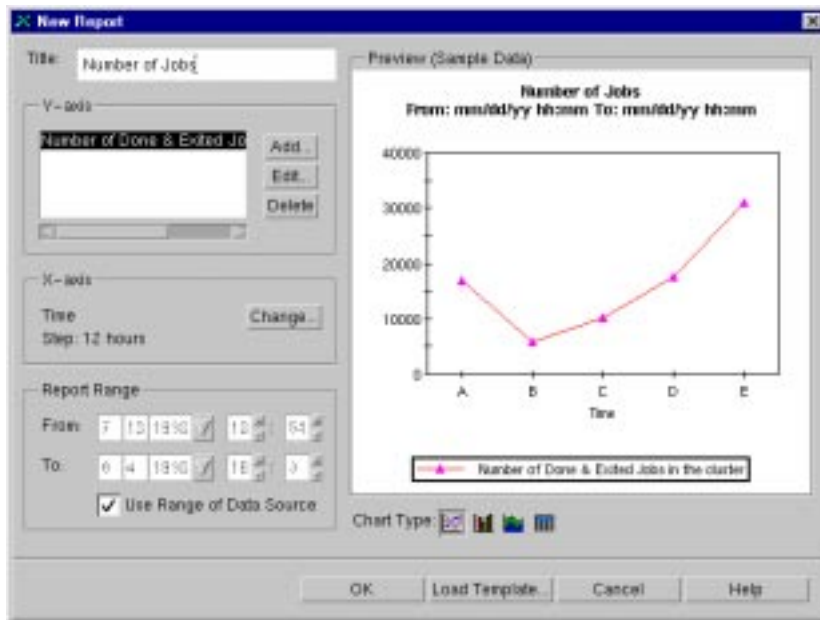
Step 7 Click **OK**, displays the *Number of Jobs* report.
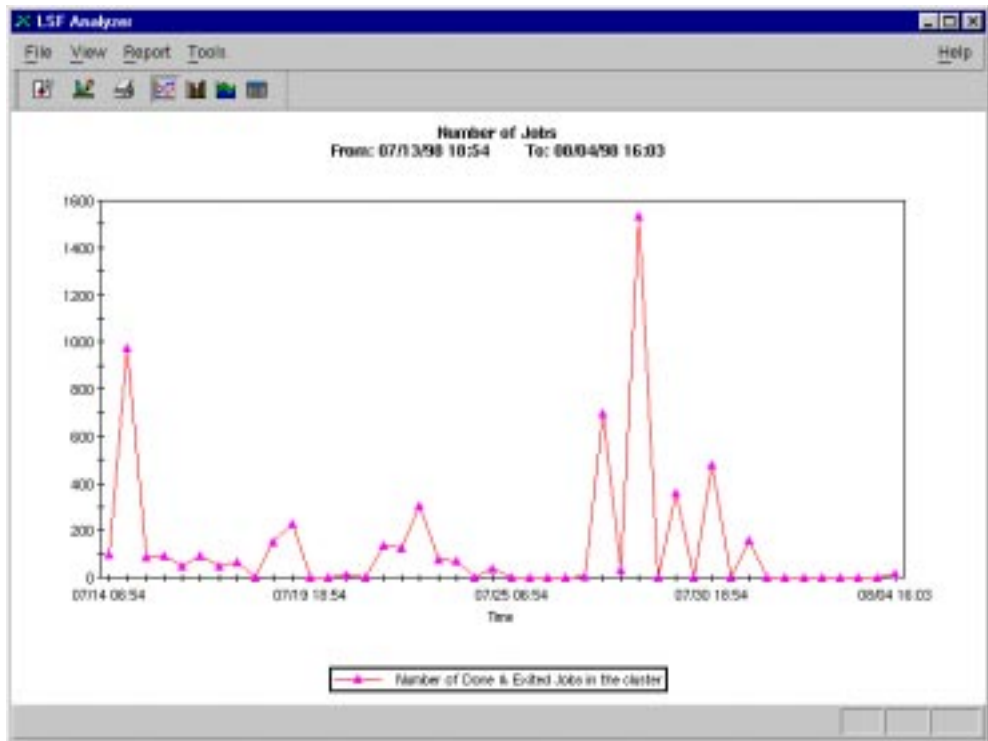
Modifying the Chart Type

To change the line chart to a bar graph, area graph or table select **Report | Bar**, **Area**, or **Table**.

2   Getting Started

3. Generating Reports

About Reports

In LSF Analyzer, a report displays one or more statistics on the Y-axis, measured against a single set of values on the X-axis.
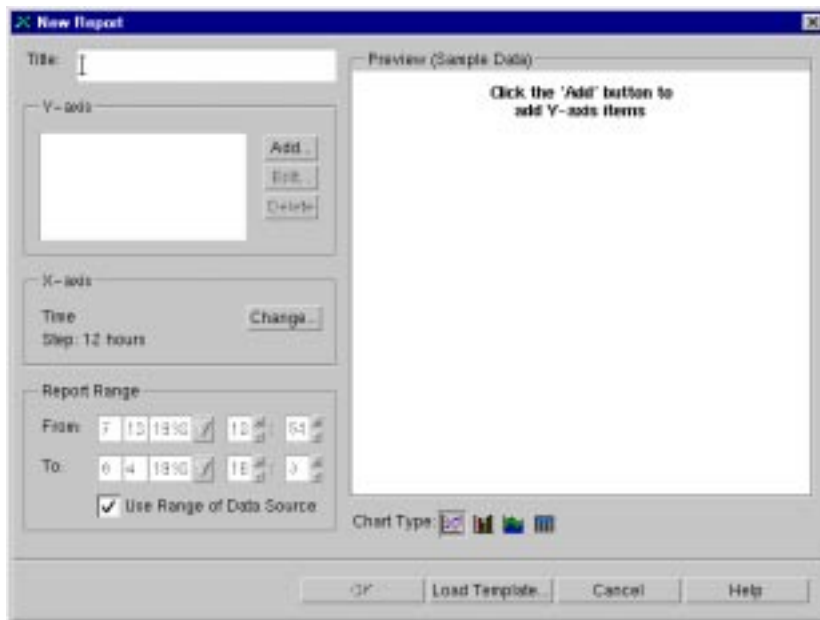
The report contains the following elements:

Data Source
> database containing the information used to generate the statistics (the data source is specified when starting LSF Analyzer but is not displayed in the report)

Title

user-specified title displayed on the report (optional)

Y-axis

statistics calculated by LSF Analyzer (maximum 7 per report)

X-axis

values (such as hosts, users, projects, or time periods) for which the statistics are calculated

Report Range

the period of time over which the statistics are calculated

Chart Type

the method of displaying the information (line chart, bar chart, area chart, or table format)

Selecting the Data Source

Each LSF cluster writes cluster information to a database, and this information is used to generate the LSF Analyzer reports. You will have multiple databases if old databases have been archived, or if you have multiple clusters. Before you can create a report, you must specify which database contains the information you want to analyze. By default, the database uses the same name as your cluster (e.g., `cluster1`).

To specify a database, take the following steps.

Step 1 Choose **File | Select Data Source**.

Step 2 In the Select Data Source dialog, select one data source, then click OK.

Y-axis Statistics

The statistics calculated by LSF Analyzer are chosen in the Y-axis dialog. A report can display up to seven statistics against suitable values of the X-axis.
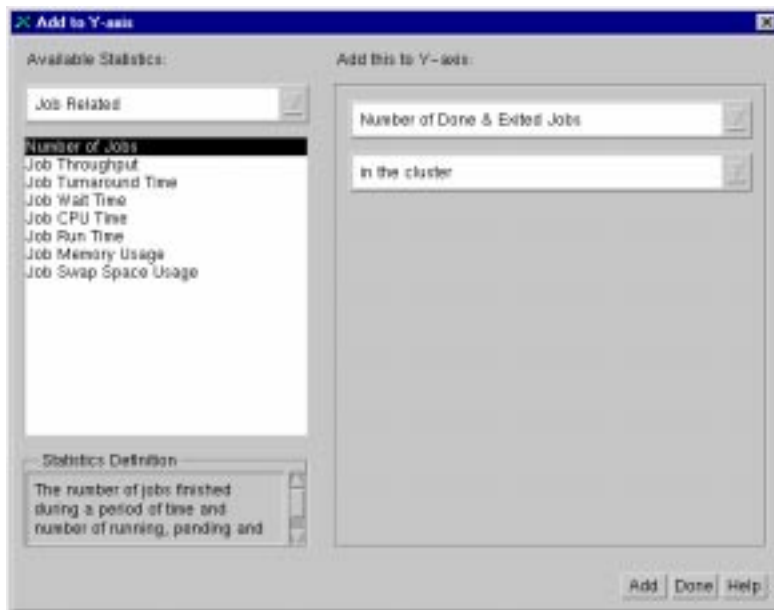
There are two ways of interpreting Time on the X-axis, depending on the Y-axis statistics. *Accumulated* statistics use data accumulated over a user-specified time interval, and report the total or an average value. *Sampled* statistics use data that is sampled at a single point in time, and the data sampling is repeated at user-specified time intervals.

The statistics are sorted by categories and classes in the Y-axis dialog, which are described in detail in *Appendix A, 'Categories of Statistics',* on page 59 and *Appendix B, 'Classes of Statistics',* on page 61.

These tables show all the statistics available on the Y-axis. They are also listed alphabetically in *Appendix C, 'Statistics',* on page 65.

3 Generating Reports

The following statistics can be displayed over these X-axis values: Time (accumulated), Hosts, Host Types, Host Models, Queues, Users, Projects, Job Names, Job CPU Time Ranges, Job Memory Usage Ranges, and Job Swap Space Usage Ranges.

Statistic	Description
Number of Done and Exited Jobs	All finished jobs (includes both successful and failed jobs, but does not include running, pending, or suspended jobs).
Number of Done Jobs	Successful jobs (jobs that were successfully dispatched and completed without an error status).
Number of Exited Jobs	Failed jobs (jobs that were not successfully dispatched or exited with an error status).

The following statistics can be displayed over these X-axis values: Time (accumulated), Hosts, Host Types, Host Models, and Queues.

Statistic	Description
Job Throughput	The number of done and exited jobs divided by the time taken to finish these jobs (in jobs/hour).

The following statistics can be displayed over these X-axis values: Time (accumulated), Queues, Users, Projects, and Job Names.

Statistic	Description
Average Turnaround Time of Done & Exited Jobs	Average time from submission to completion for all finished jobs (in seconds).
Average Turnaround Time of Done Jobs	Average time from submission to completion for successful jobs (in seconds).
Average Turnaround Time of Exited Jobs	Average time from submission to completion for failed jobs (in seconds).

Statistic	Description
Average Wait Time of Done & Exited Jobs	Average time from submission to dispatch for all finished jobs (in seconds).
Average Wait Time of Done Jobs	Average time from submission to dispatch for successful jobs (in seconds).
Average Wait Time of Exited Jobs	Average time from submission to dispatch for failed jobs (in seconds).

The following statistics can be displayed over these X-axis values: Time (accumulated), Hosts, Host Types, Host Models, Queues, Users, Projects, and Job Names.

Statistics	Description
Average CPU Time of Done & Exited Jobs	Average time the CPU spent running each job in user mode and in kernel mode (in seconds).
Average CPU Time of Done Jobs	Average time the CPU spent running each successful job in user mode and in kernel mode (in seconds).
Average CPU Time of Exited Jobs	Average time the CPU spent running each failed job in user mode and in kernel mode (in seconds).
Total CPU Time of Done & Exited Jobs	Total CPU time spent running all jobs in user mode and in kernel mode (in seconds).
Total CPU Time of Done Jobs	Total CPU time spent running successful jobs in user mode and in kernel mode (in seconds).
Total CPU Time of Exited Jobs	Total CPU time spent running failed jobs in user mode and in kernel mode (in seconds).
Average Run Time of Done & Exited Jobs	Average time from dispatch to completion for all finished jobs (in seconds).
Average Run Time of Done Jobs	Average time from dispatch to completion for successful jobs (in seconds).
Average Run Time of Exited Jobs	Average time from dispatch to completion for failed jobs (in seconds).

Statistics	Description
Total Run Time of Done & Exited Jobs	Time from dispatch to completion, total for all finished jobs (in seconds).
Total Run Time of Done Jobs	Time from dispatch to completion, total for all successful jobs (in seconds).
Total Run Time of Exited Jobs	Time from dispatch to completion, total for all failed jobs (in seconds).
Average Memory Usage of Done & Exited Jobs	Average maximum resident memory used by processes of all finished jobs (in KB).
Average Memory Usage of Done Jobs	Average maximum resident memory used by processes of successful jobs (in KB).
Average Memory Usage of Exited Jobs	Average maximum resident memory used by processes of failed jobs (in KB).
Total Memory Usage of Done & Exited Jobs	Maximum resident memory used by all processes of a job, total for all finished jobs (in KB).
Total Memory Usage of Done Jobs	Maximum resident memory used by all processes of a job, total for successful jobs (in KB).
Total Memory Usage of Exited Jobs	Maximum resident memory used by all processes of a job, total for failed jobs (in KB).
Average Swap Usage of Done & Exited Jobs	Average maximum virtual memory used by processes of all finished jobs (in KB).
Average Swap Usage of Done Jobs	Average maximum virtual memory used by processes of successful jobs (in KB).
Average Swap Usage of Exited Jobs	Average maximum virtual memory used by processes of failed jobs (in KB).
Total Swap Usage of Done & Exited Jobs	Maximum virtual memory used by all processes of a job, total for all finished jobs (in KB).
Total Swap Usage of Done Jobs	Maximum virtual memory used by all processes of a job, total for successful jobs (in KB).
Total Swap Usage of Exited Jobs	Maximum virtual memory used by all processes of a job, total for failed jobs (in KB).

The following statistics can only be displayed over Time (accumulated) on the X-axis

Statistic	Description
CPU Utilization	The CPU time used over the last minute divided by the CPU time available in the same period (in percentage).
Memory Utilization	The amount of memory used divided by the total amount of memory available on the host (in percentage).
Swap Space Utilization	The amount of swap space used, divided by the total amount of swap space available on the host (in percentage).
Batch Job Slot Utilization	The number of used batch job slots divided by the maximum number of job slots on the host (in percentage).
Batch Processor Utilization	The number of used batch job slots divided by the number of processors on the host (in percentage).
15-second Run Queue Length	The 15-second exponentially averaged CPU run queue length.
1-minute Run Queue Length	The 1-minute exponentially averaged CPU run queue length.
15-minute Run Queue Length	The 15-minute exponentially averaged CPU run queue length.
Paging Rate	The memory paging rate exponentially averaged over the last minute (in pages/second).
Disk I/O Rate	The disk I/O rate exponentially averaged over the last minute (in KB/second).
Number of Login Users	The number of current login users.
Interactive Idle Time	For all logged in sessions, the amount of time during which the keyboard is not used (in minutes).
Available Memory	Amount of memory available (in MB).
Available Swap Space	Amount of swap space available (in MB).

Statistic	Description
Available /tmp Space	Amount of space available in the temporary directory (in MB).
User-specified external load indices	Any user-defined external dynamic numeric resource.
Usage of Resource Shared among Hosts	Total value of all instances for a dynamic shared resource.

The following statistics can only be displayed over Time (sampled) on the X-axis.

Statistic	Description
Num of Running, Pending & Suspended Jobs	Jobs in the system, either running or waiting to run
Num of Running & Suspended Jobs	Jobs which have been dispatched but not finished (running jobs or jobs in a suspended state)
Num of Pending Jobs	Jobs submitted but not yet dispatched.
Average CPU Time of Running & Suspended Jobs	CPU time used by each job, averaged for all jobs which have been dispatched but not finished (in seconds).
Total CPU Time of Running & Suspended Jobs	Total time the CPU has spent running jobs which have been dispatched but not finished, in user mode and in kernel mode (in seconds).
Average Memory Usage of Running & Suspended Jobs	Resident memory used by all processes in a job, averaged for all jobs which have been dispatched but not finished (in KB).
Total Memory Usage of Running & Suspended Jobs	Total resident memory used by all processes in jobs which have been dispatched but not finished (in KB).
Average Swap Usage of Running & Suspended Jobs	Virtual memory used by all processes in a job, averaged for all jobs which have been dispatched but not finished (in KB).
Total Swap Usage of Running & Suspended Jobs	Total virtual memory used by all processes in jobs which have been dispatched but not finished (in KB).

Statistic	Description
Number of Hosts in the Cluster	Number of LSF batch server hosts in the cluster.
Number of OK Hosts	Hosts able to accept batch jobs
Number of Busy Hosts	Overloaded hosts which are unable to accept batch jobs because some load indices go beyond the configured thresholds.
Number of Full Hosts	Hosts which are unable to accept batch jobs because the configured maximum number of batch job slots has been reached.
Number of Closed Hosts	Hosts which are unable to accept batch jobs for any of the following reasons: they are running an exclusive job, they have been locked by the LSF administrator, or they have been closed by the LSF administrator or its dispatch windows.
Number of Unavailable Hosts	Hosts which are unable to accept batch jobs because they are down, or their LIM/sbatchd is unreachable.

X-axis Values

The statistics you display on the Y-axis are measured against values on the X-axis.

Statistics can be displayed against the following values:

- Time

- Hosts

- Host Types

- Host Models

- Queues

- Users

- Projects

- Job Names

- Job CPU Time Ranges

- Job Memory Usage Ranges

- Job Swap Space Usage Ranges

Report Range

You may choose to limit the report range so the report is generated faster, and so you can easily analyze data collected over a time period that is meaningful to you. You can specify the time and date of the start and end of the report range.

Accumulated statistics are calculated using data written to the database between the times and dates specified. By default, the statistics will be calculated using all of the data in the database.

Sampled statistics are displayed at regular intervals (specified in the Y-axis dialog) between the times and dates specified. By default, these statistics are displayed over the entire range of the data.

Creating Reports

When creating a report, you can make changes to all the elements in the New Report dialog, until the report displays the information you want. The report is generated when you click OK, and the time it takes depends on the amount of data involved.

To create a new report, take the following steps:

Step 1 Choose **File | New Report**.

Step 2 Specify the title of the report.

Step 3   Click Add to specify the first statistic you want to display on the Y-axis. In the Y-axis dialog, select the category and class on the left side, then specify the statistic and scope on the right side.

Click Edit or Delete if you want to modify any of the Y-axis statistics you have selected.

You may include up to seven statistics in one report.

Step 4   Specify the values of the X-axis.

Step 5   Specify the Report Range. By default, the entire range of the database is used.

Step 6   Specify the Chart Type (line chart, bar chart, area chart, or table format).

Step 7   Click OK.

It is easier to create a report from an existing template than it is to create a report from scratch. For more information, see *'Using a Template to Create a Report'* on page 29.

Modifying Reports

To modify the current report, take the following steps:

Step 1   Click the Modify button on the toolbar, or choose **Report | Modify**.

Step 2   In the Modify Report dialog, make any changes you want, then click OK.

For more information, see *'Using a Template to Modify a Report'* on page 30.

Printing Reports

To print the current report, take the following steps:

Step 1   Choose **File | Print**.

Step 2 In the Print dialog, specify printing parameters, then click OK.

Exporting Reports

To export data from the current report into a CSV (comma-separated) file, take the following steps:

Step 1 Choose **File | Export**.

Step 2 In the File Selection dialog, type or select the CSV file name, then click OK or Save.

3 Generating Reports

4. Using Templates

About Templates

LSF Analyzer includes templates to help you reproduce reports. A template stores all the characteristics of an analysis report (title, Y-axis, X-axis, report range and chart type) except for the data itself.

You can save any report as a template and use it as the basis for generating a similar analysis of different job data.

You can use any template to create a new report or modify an existing one.

Saving a Template

To save the current report as a template, take the following steps:

Step 1   Choose **File | Save as Template**.

Step 2   In the File Selection dialog, type or select the template file name, and click OK or Save.

Using a Template to Create a Report

To create a report from a template, take the following steps:

Step 1   Choose **File | New Report from Template**.

Step 2 In the Load Template dialog, type or select the template file name, and click OK or Open.

Step 3 In the New Report dialog, make any changes you want, then click OK.

Using a Template to Modify a Report

Using a template to modify a report is the same as deleting the report and creating a new report from a template.

To modify the current report using a template, take the following steps:

Step 1 Click the Modify button on the toolbar, or choose **Report | Modify**.

Step 2 In the Modify Report dialog, click Load Template.

Step 3 In the File Selection dialog, type or select the template file name, and click OK or Open.

Step 4 In the Modify Report dialog, make any changes you want, then click OK.

4

Default Templates

A wide variety of default templates can be installed with LSF Analyzer and used to create reports quickly. They are found in the `%LSF_MISC%\xanalyzer` directory on Windows NT, and under the `$LSF_MISC/Xanalyzer` directory on UNIX.

The templates are organized as shown:

Accounting
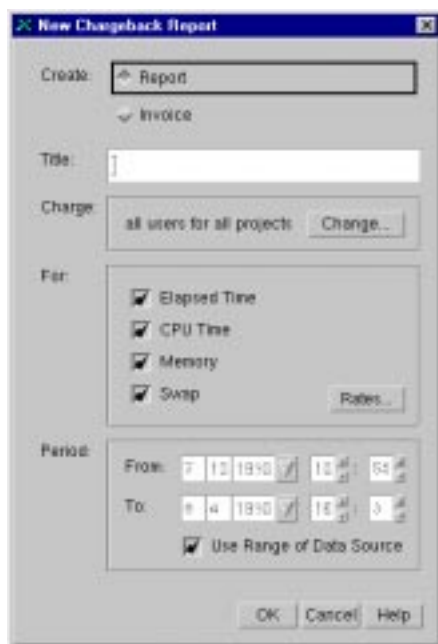>Project
>Queue
>User

Performance
>General
>Host
>Queue

Workload
>General
>Job_Profile
>Job_Type

4 Using Templates

5. Chargeback Accounting

LSF Analyzer can be used to perform chargeback accounting by generating tabular reports and invoices. The chargeback accounting uses the same statistics as LSF Analyzer reports, but a cost is associated with the use of each resource. You can determine the costs associated with a set of users or projects, over a specified time range.

About Chargeback Reports

5 Chargeback Accounting

Chargeback reports contain all of the following elements:

Title
> user-specified title displayed on the report (optional)

Format
> method of displaying the information (Report or Invoice format)

Who to Charge
> who or what will be charged for resource usage (Users or Projects)

Rates
> costs associated with the use of each resource

Resources to Charge For
> one or more resources to account for in the report
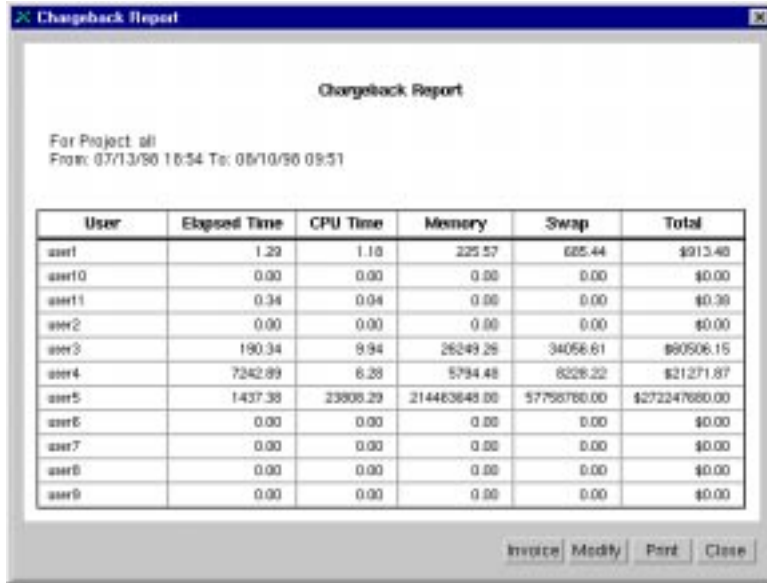
Report Range
> the period of time over which the charges are calculated

5

Reports and Invoices

A chargeback report shows the costs of resources used by individual users or projects, and also the total costs of resources used, in a table format.

5 Chargeback Accounting

A set of invoices is just a different way of displaying the same information contained in a report. An invoice gives the cost of resources used by individual users or projects. One invoice is generated for each user or project included in the report.

Invoices may be used for billing purposes. The tabular format may be used for comparison purposes, and to summarize and archive periodic reports.

5
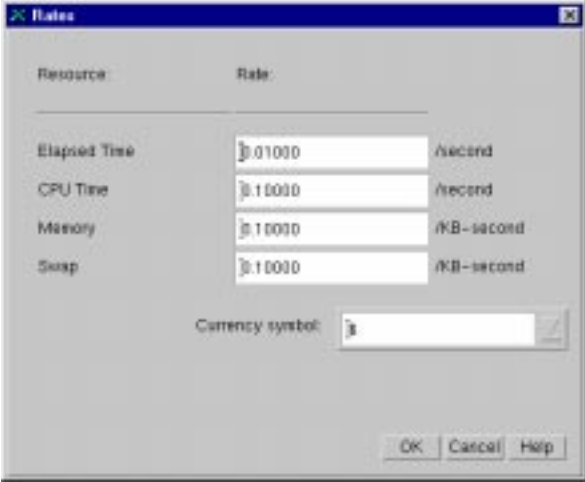
Chargeback Resources

LSF Analyzer can generate chargeback reports for any combination of resources. The resources for which chargeback accounting is done are listed below.

Resource	Description	Unit
Elapsed Time	The wall clock time from the start of a job to its end. This is useful if you want to charge for the actual time a job is running in the cluster. This time includes any suspension of a job by the person who submitted the job.	seconds
CPU Time	The CPU time used by the job. This is useful if you want to charge for the actual use of CPU resources by a job. Computation-intensive jobs will be charged more than other types of jobs.	seconds
Memory	The maximum resident memory used by the job, multiplied by the elapsed time. This is useful if you want to charge for the maximum resident memory used by the job during its run time.	KB-seconds
Swap	The maximum virtual memory used for the job, multiplied by the elapsed time. This is useful if you want to charge for the total size of the core images used by a job during its run time.	KB-seconds

Chargeback Rates

By default, LSF Analyzer gets chargeback rates from the `lsb.charge.rate` configuration file, located in the `LSF_MISC` directory of the cluster. You can easily view or change the rates by clicking the Rates button in the New Chargeback Report dialog or Modify Chargeback Report dialog.

This basic formula is used to calculate the cost of a particular resource:

        number of units used * resource charge rate

When you set the rates, you must specify the currency. The charge rate can be specified for each separate resource. The unit of measurement associated with each resource cannot be changed.

Generating Chargeback Reports

To create a new chargeback report, take the following steps:

Step 1   Choose **Tools | Chargeback**.

Step 2 In the New Chargeback Report dialog, specify the following elements:

- title (optional)

- format (Report or Invoice)

- who to charge (Users or Projects)

- what resource usage to charge for

- resource rates (Rates button; optional)

- report range

Step 3 Click OK.

Modifying Chargeback Reports

To modify the current chargeback report or invoice, take the following steps:

Step 1 In the Chargeback Report dialog, click the Modify button.

Step 2 In the Modify Chargeback Report dialog, make any changes you want.

Step 3 Optional. If you want to change the rates charged for resources, click the Rates button. Make changes to the Rates dialog, then click OK

Step 4 Click OK in the Modify Chargeback Report dialog.

Printing Chargeback Reports

To print the current chargeback report, take the following steps:

Step 1 Click **Print** on the Report dialog.

Step 2 In the Print dialog, specify printing parameters, then click **OK**.

Printing Chargeback Invoices

To print chargeback invoices, take the following steps:

Step 1 Select a user name from the User list.

Step 2 Click **Print** on the Invoice dialog.

Step 3 In the Print dialog, specify printing parameters, then click OK.

6. LSF Database - UNIX

This chapter is aimed at the LSF administrator running a UNIX cluster or a mixed UNIX/NT cluster. If your cluster contains only Windows NT hosts, see *Chapter 7, 'LSF Database - Windows NT', on page 49.*

Note

> *LSF database administration is the only part of LSF that requires you to log on using the LSF primary administrator user account. The LSF primary administrator account will be the first cluster administrator specified in the* `lsf.cluster.cluster` *file. If your cluster includes any Windows NT hosts, do not attempt to change the LSF primary administrator after LSF has been installed.*

In This Chapter

LSF Database Installation

Installation

Install LSF Analyzer using `lsfsetup`, the UNIX installation program for LSF Suite. For detailed instructions, see the *LSF Installation Guide.*

When you install LSF Analyzer for the first time, you are prompted to specify the LSF database server. This defines the `LSF_DB_HOST` parameter in the `lsf.conf` file. The LSF database files will be installed on the machine you specify.

- If you specify the installation host (the machine used to run `lsfsetup`) as the LSF database server, the LSF database installation program starts automatically.

- If you specify another machine as the LSF database server, you have to install the LSF database using a separate installation file that is included in the regular LSF distribution. Complete the regular LSF installation, using `lsfsetup,` and then log onto the LSF database server and run `lsfdbsetup.`

When you install the LSF database, you are prompted to specify the LSF database directory. By default, this is `/usr/local/lsf_db`. All the LSF database files are installed under this directory. The database is created automatically and has the same name as the cluster.

Starting the Database

Once the database is installed, you must start it.

To start the LSF database, log onto the LSF database server as the LSF primary administrator and type:

```
% lsdbserver start
```

The `lsdbserver` command is located in the `/bin` subdirectory of the LSF database directory.

LSF Database Utility Commands

The utilities described in this section help you monitor and limit the size of the LSF database, to get better performance from LSF Analyzer. They are used to:

- start up and shut down the LSF database

- create and drop LSF databases

- move or delete existing data in the LSF database

- show the status of the LSF database

- show the size of the LSF database

- rebuild the indices of the LSF database

Environment Configuration

The `lsf.conf` file and the online documentation for database utility commands should be accessible to the working environment of the LSF primary administrator on the LSF database server (`LSF_ENVDIR` must be set). For example:

```
% setenv PATH $DB/bin:$PATH
% setenv MANPATH $DB/man:$MANPATH
```

where `$DB` is the LSF database directory.

Command Reference

To run any of the commands in this section, you must log onto the LSF database server as the LSF primary administrator. The following commands all support the standard `-h` option.

lsdbserver start|stop
> Starts up or shuts down the LSF database.

lsdbcreate *database_name*

Creates a new database with the LSF database schema. Specify a unique database name (you cannot use the name of an existing LSF database).

lsdbstatus

Prints the current status of the LSF database server (which configuration file is in use; maximum and actual number of connections; details of all connections).

lsdbrecords *database_name* **[-t** *time* **]**

Prints reference numbers for job-related and resource-related records. If –t is specified, prints another pair of reference numbers, for job and resource records logged before the specified time.

If you want good performance from LSF Analyzer, you should keep the LSF database small. When the reference numbers exceed 50,000 for job records or 400,000 for resource records, you should archive the data using the lsdmove command.

The rate of growth of your LSF database depends on all of the following:

- number of jobs issued daily in the cluster

- number of hosts in the cluster

- number of resource indices logged for the cluster

-t *time*

To specify time, use the following syntax:

```
time = [year /] month / day / hour : minute
```

Use 4 digits for the year, if specified. Use 2 digits for month, day, hour, and minute.

lsdbbuildidx *database_name*

Reindexes the specified LSF database. If LSF Analyzer users have trouble retrieving data from the database, try this command, as there might be a problem with the database index.

lsdbclear *database_name*

Deletes all the logged data from the data tables in the specified LSF database, leaving an empty LSF database.

CAUTION!

There is no way to recover the deleted data.

Note

You should not run this command on a working database, such as the online LSF database. If you do, use the lsdbstatus *command and make sure there are no open connections to the database before using this command. This will prevent data inconsistency problems.*

lsdbdrop *database_name*

Drops the specified LSF database (deletes the entire LSF database, including all of the logged data).

CAUTION!

There is no way to recover a dropped LSF database.

lsdbmove *source_database_name destination_database_name*

Moves the contents of the source database (data) to the destination database. This command must be followed by the badmin reconfig command.

LSF Data Collection Parameters

Data collection can be tuned by modifying the parameters configured in the lsb.params file. For example, edit:

> *LSB_CONFDIR/cluster*/configdir/lsb.params

The following parameters can be configured:

DB_DEFAULT_INTVAL

Specifies the time interval to log job data (except load information and resource usage of running and suspended jobs) to the database, in minutes. To stop logging job data, set DB_DEFAULT_INTVAL=-1.

Default: 5

DB_JOB_RES_USAGE_INTVAL

> Optional. If defined, specifies the time interval to log resource usage of running and suspended jobs to the database, in minutes.

> Default: undefined (job resource is not logged).

DB_LOAD_INTVAL

> Specifies the time interval to log load information (internal and external load indices and shared resources), in minutes.

> Minimum: 15

> Default: 60 (1 hour)

DB_SELECT_LOAD

> Optional. Specifies which load information to collect. Possible values are internal load indices, external load indices, and shared resources. This parameter is case sensitive. If more than 4 load values are specified (separated by spaces) only the first 4 will be used. For example,

> ```
> DB_SELECT_LOAD = ut mem ext_idx1 shared_lic1
> ```

> will collect CPU utilization, available memory, the user-specified external load index `ext_idx1`, and the user-specified shared resource `shared_lic1`.

> Default: `ut`

lsb.acct Data Conversion

The `acct2db` utility is used to convert job log files (e.g., `lsb.acct`) into LSF databases, allowing you to analyze LSF data collected before LSF Analyzer was installed. These databases cannot be used as online (active) databases. The license for the `acct2db` utility expires 30 days after your LSF 3.2 license is generated.

To convert data from an existing job log file, take the following steps:

Step 1 Use the LSF primary administrator user account to log onto the host where the job log file (e.g., `lsb.acct`) is stored.

Step 2 Create a new LSF database using the `lsdbcreate` command described in the *'Command Reference'* on page 44.

Step 3 Issue the `acct2db` command and specify the new database as the target database. Each finished job record in the job log file is converted into a database record.

Syntax

```
acct2db [ -h ] [ -V ]
        [ -f lsb.acct.file ]
        [-H database_host ]
        database_name
```

-h

> Prints command usage to stderr and exits.

-V

> Prints the LSF release version to stderr and exits.

-f *lsb.acct.file*

> Specifies the job log file on the local host which is to be converted into a database.
>
> Default: LSB_SHAREDIR/*cluster*/*logdir*/lsb.acct

-H *database_host*
> Specifies the remote host where the target database is located.
>
> Default: the local host

database_name
> Specifies the name of the target LSF database. The target database cannot be the online (active) database.

7. LSF Database - Windows NT

This chapter is aimed at the LSF administrator running a Windows NT only cluster. If your cluster contains any UNIX hosts, see *Chapter 6, 'LSF Database - UNIX', on page 41.*

All examples in this chapter assume the use of Microsoft SQL Server 6.5 as the LSF database, and use of an ODBC-compliant DBMS.

In This Chapter

Database Installation Procedures

- *'Database Setup'* on page 50

- *'Host Setup'* on page 51

- *'LSF Setup'* on page 53

Database Commands and Parameters

- *'LSF Database Utility Command'* on page 56

- *'LSF Data Collection Parameters'* on page 56

- *'lsb.acct Data Conversion'* on page 57

Database Setup

Log on as the Microsoft SQL Server Database Administrator (DBA) and complete the following procedures in order.

*   *'Create LSF Database Accounts'* on page 50

*   *'Create A New Database'* on page 51

*   *'Build LSF Database Schema'* on page 51

*   *'Grant Permissions to LSF Database Accounts'* on page 51

Create LSF Database Accounts

Step 1   Create a new account that will be used by LSF Batch to write information to the LSF database. This is the LSF database user account.

By default, LSF Analyzer also uses this account to read information in the LSF database.

Step 2   If you are very concerned about security, you may choose to create a second account that does not have permission to write to the database. This is the optional LSF database guest account.

LSF Analyzer does not require permission to write to the database, so if you create the guest account, LSF Analyzer will use it instead of the LSF database user account. The guest account could also be used by other users who only need to read the contents of the database.

Note

To simplify database setup, give the LSF database user account the same name as your LSF primary administrator user account, and do not create the LSF database guest account. You can specify any password you want, but do not confuse the database password with the password for the Windows NT user account.

Create A New Database

Create a database that will be the new LSF database.

Note

> *To simplify database setup, give the database the same name as your LSF cluster (e.g.,* `cluster1`*).*

Build LSF Database Schema

Create the tables in the new database according to the LSF database schema. The LSF database schema is specified in a Microsoft SQL Server 6.5 ISQL executable file named `createschema.sql`, located in the *cluster*\bin directory (e.g., cluster1\bin).

Grant Permissions to LSF Database Accounts

Grant the LSF database user account Select, Delete, and Update permissions for all tables in the new database.

If you created an LSF database guest account, grant it Read permission for the new database.

Host Setup

Before you can use LSF Analyzer, you must install the ODBC driver and set up a Data Source Name (DSN) on each batch host in the cluster and on the host used to run the LSF Analyzer (xanalyzer) graphical user interface.

Log onto each host and, if necessary, install the ODBC driver before you set up the DSN.

ODBC Driver Installation

LSF Analyzer requires the version 3.0 or later ODBC driver for Microsoft SQL Server 6.5. The driver is available from Microsoft.

Step 1 Download the Microsoft ODBC 3.5 SDK (ODBC35IN.exe) which installs the supported ODBC driver for Microsoft SQL Server 6.5 from:

```
http://www.microsoft.com/data/odbc
```

Step 2 Run ODBC35IN.exe to install the Setup for Microsoft ODBC 3.5 SDK

Step 3 Run Setup.exe from the ODBC 3.5 SDK install directory to begin the ODBC driver installation. Select the Minimum installation option.

DSN Setup

Step 1 Run **ODBC Administrator** from the Windows Control Panel.

Step 2 Select the **System DSN** tab.

Step 3 Click the **Add** button. Displays the **Create New Data Source** dialog.

Step 4 Select the **MS SQL** driver from the list. Displays the **Create New Data Source for SQL Server** dialog.

Step 5 Enter the name of your MS-SQL database in the **Name** field and add the **.ana** extension (e.g., cluster1.ana).

Step 6 Select the LSF database server from the **Server** list.

Step 7 Enter a description for the DSN into the **Description** field. Click **Next**.

Step 8 Select the **With SQL Server authentication using a login ID & Password entered by the user** option. Enables the **Login ID** and **Password** fields.

Step 9 Enter the name of the LSF database account for **Login ID** (e.g., lsfadmin) and the password of the LSF database account for **Password**. Click **Next**.

Step 10 Select the **Change the default database to** option and enter the name of the
LSF database (e.g., `cluster1`). Click **Next**.

Step 11 Follow the prompts and accept all default values to complete the
configuration. Click **Finish** when displayed. Displays the **Confirmation**
dialog.

Step 12 Click **Test**.

Step 13 After the test completes, click **OK**.

LSF Setup

Log on as an LSF cluster administrator (not the LSF primary administrator).

The following process describes the activities for the LSF cluster administrator to carry
out for setting up LSF for use with Microsoft SQL Server 6.5 and LSF Analyzer:

*   *'Set Database Login Passwords'* on page 53

*   *'Set LSF Database Parameters'* on page 54

*   *'Update LSF'* on page 54

Set Database Login Passwords

LSF needs to know the database password for the LSF database user account and the
LSF database guest account, if it was created (see *'Create LSF Database Accounts'* on
page 50).

Use the `lsdbpasswd` command to set the passwords in LSF, see *'lsdbpasswd'* on
page 56. This command will ask for the password for the specified user ID, encrypt the
password and save it into a file. For example, type:

```
C:> lsdbpasswd lsfadmin
```

If the LSF database user account has the same name as a Windows NT user account, use this command to input the database password, not the Windows NT password.

Set LSF Database Parameters

Edit the `lsf.conf` file and set the following parameters to define the LSF database environment.

Note
> *If you took all the recommended steps to simplify database setup, you do not need to set these parameters. The default values will match your database setup.*

LSF_DB_ACTIVE
> Specifies the name of the active LSF database, see *'Step 5'* on page 52.
>
> Default: the name of the cluster is used (e.g., `cluster1`), but the variable remains undefined.

LSF_DB_ACCT
> Specifies the Login ID for the LSF database account, see *'Create LSF Database Accounts'* on page 50.
>
> Default: the name of the LSF primary administrator account is used (e.g., `lsfadmin`), but the variable remains undefined.

LSF_DB_GUEST_ACCT
> Specifies the Login ID for the LSF database guest account, see *'Create LSF Database Accounts'* on page 50.
>
> Default: the name of the LSF primary administrator account is used (e.g., `lsfadmin`), but the variable remains undefined.

Update LSF

The procedure to update LSF differs depending on whether LSF is running or yet to be started. Is LSF running?

- **Yes** - Go to Reconfigure LSF

- **No** - Go to Starting LSF

Reconfigure LSF

Reconfigure LSF by executing the following two commands:

```
C:> lsadmin reconfig
C:> badmin reconfig
```

Starting LSF

The LSF service and daemons on each LSF server host will start automatically when the machine is restarted.

If you cannot restart each host at this time, log on as an LSF cluster administrator (a member of the LSF Global Administrators group) and start the LSF service and daemons manually.

Note

> *You should not use the primary LSF administrator's account (normally* lsfadmin*) to start or stop LSF service and daemons.*

To start the LSF service and daemons, use any one of the following methods:

- Use the Windows NT Server Manager to start "LSF Service" on all LSF server hosts.

- Click "Services" on the Windows NT Control Panel and start "LSF Service". You will have to repeat this step on each LSF server host.

- If LSF Batch has been installed, go to the "LSF Suite for Workload Management/ LSF Batch" program folder, and use the LSF administrative tool "LSF Batch Administration". (You can use this tool to perform all your administrative tasks for LSF Base and LSF Batch products.)

- Start a new command console, and type:

  ```
  C:> lssrvcntrl start -m all lssrvman
  ```

Usage information for lssrvcntrl is available by typing lssrvcntrl with no
options.

LSF Database Utility Command

lsdbpasswd

lsdbpasswd [-h] *userID*
> Sets and changes the user's password. The password is encrypted then written
> to the lsf.dbpasswd file.

-h
> Print command usage to stderr and exit.

userID
> Specifies user ID.

LSF Data Collection Parameters

Data collection can be tuned by modifying the parameters configured in the
lsb.params file. For example, edit:

> *LSB_CONFDIR**cluster*\configdir\lsb.params

The following parameters can be configured:

DB_DEFAULT_INTVAL
> Specifies the time interval to log job data (except load information and
> resource usage of running and suspended jobs) to the database, in minutes. To
> stop logging job data, set DB_DEFAULT_INTVAL=-1.

> Default: 5

`DB_JOB_RES_USAGE_INTVAL`

> Optional. If defined, specifies the time interval to log job resource usage to the database, in minutes.
>
> Default: undefined (job resource is not logged).

`DB_LOAD_INTVAL`

> Specifies the time interval to log load information (internal and external load indices and shared resources), in minutes.
>
> Minimum: 15
>
> Default: 60 (1 hour)

`DB_SELECT_LOAD`

> Optional. Specifies which load information to collect. Possible values internal load indices, external load indices, and shared resources. This parameter is case sensitive. If more than 4 load values are specified (separated by spaces) only the first 4 will be used. For example,
>
> ```
> DB_SELECT_LOAD = ut mem ext_idx1 shared_lic1
> ```
>
> will collect CPU utilization, available memory, the user-specified external load index `ext_idx1`, and the user-specified shared resource `shared_lic1`.
>
> Default: `ut`

lsb.acct Data Conversion

The `acct2db` utility is used to convert job log files (e.g., `lsb.acct`) into LSF databases, allowing you to analyze LSF data collected before LSF Analyzer was installed. These databases cannot be used as online (active) databases. The license for the `acct2db` utility expires 30 days after your LSF 3.2 license is generated.

To convert data from an existing job log file, take the following steps (all Windows NT examples assume the use of Microsoft SQL Server 6.5 as the LSF database, and use of an ODBC-compliant DBMS). Log on as the Microsoft SQL Server Database Administrator to complete steps 1 through 4.

Step 1 Create a new LSF database, see *'Create A New Database'* on page 51.

Step 2 Build the schema, see *'Build LSF Database Schema'* on page 51.

Step 3 Grant the LSF database user account permission to read and write to the new database. If you have created an LSF database guest account, grant it Read permission for the new database.

Step 4 Set up a DSN for the new database on every batch host, and on the host where LSF Analyzer is installed. See *'DSN Setup'* on page 52.

Step 5 Log onto the host where the job log file (i.e., `lsb.acct`) is stored.

Step 6 Issue the `acct2db` command and specify the new database as the target database. Each finished job record in the job log file is converted into a database record.

Syntax

`acct2db [ -h ] [ -V ] [ -f `*`lsb.acct.file`*` ] `*`database_name`*

`-h`

> Prints command usage to stderr and exits.

`-V`

> Prints the LSF release version to stderr and exits.

`-f `*`lsb.acct.file`*

> Specifies the job log file on the local host which is to be converted into a database.

> Default: `LSB_SHAREDIR`*`cluster`**`logdir`*\\`lsb.acct`

`database_name`

> Specifies the name of the target LSF database. The target database cannot be the online (active) database.

A. Categories of Statistics

This list shows the classes of statistics (second field in the Y-axis dialog) available in each category (first field in the Y-axis dialog).

Job-related statistics
> Number of Jobs
> Job Throughput
> Job Turnaround Time
> Job Wait Time
> Job CPU Time
> Job Run Time
> Job Memory Usage
> Job Swap Space Usage

Resource-related statistics
> Usage of Resource Shared among Hosts
> Load Index of Hosts
> CPU Utilization of Hosts
> Memory Utilization of Hosts
> Swap Space Utilization of Hosts
> Batch Job Slot Utilization of Hosts
> Job CPU Time
> Job Memory Usage
> Job Swap Space Usage

Host-related statistics
> Number of Hosts
> Load Index of Hosts
> CPU Utilization of Hosts
> Memory Utilization of Hosts
> Swap Space Utilization of Hosts
> Batch Job Slot Utilization of Hosts
> Number of Jobs

A Categories of Statistics

Job Throughput
Job Turnaround Time

Queue-related statistics
Number of Jobs
Job Throughput
Job Turnaround Time
Job Wait Time
Job CPU Time
Job Run Time
Job Memory Usage
Job Swap Space Usage

User-related statistics
Number of Jobs
Job Turnaround Time
Job Wait Time
Job CPU Time
Job Run Time
Job Memory Usage
Job Swap Space Usage

Project-related statistics
Number of Jobs
Job Turnaround Time
Job Wait Time
Job CPU Time
Job Run Time
Job Memory Usage
Job Swap Space Usage

B. Classes of Statistics

This list shows the statistics (third field in the Y-axis dialog) in each class (second field on the Y-axis dialog). For each class, the definition and statistics are shown.

Batch Job Slot Utilization of Hosts

Batch Job Slot Utilization
Batch Processor Utilization
Definitions: The number of used batch job slots divided by the maximum number of job slots on the host and by the number of processors on the host (in percentage).

CPU Utilization of Hosts

CPU Utilization
Definition: The CPU time used over the last minute divided by the CPU time available in the same period (in percentage).

Job CPU Time

Average CPU Time of Done & Exited Jobs
Average CPU Time of Done Jobs
Average CPU Time of Exited Jobs
Average CPU Time of Running & Suspended Jobs
Total CPU Time of Done & Exited Jobs
Total CPU Time of Done Jobs
Total CPU Time of Exited Jobs
Total CPU Time of Running & Suspended Jobs
Definition: The time the CPU spent running job in user mode and in kernel mode (in seconds).

Job Memory Usage

Average Memory Usage of Done & Exited Jobs
Average Memory Usage of Done Jobs
Average Memory Usage of Exited Jobs
Average Memory Usage of Running & Suspended Jobs
Total Memory Usage of Done & Exited Jobs
Total Memory Usage of Done Jobs

Total Memory Usage of Exited Jobs
Total Memory Usage of Running & Suspended Jobs
Definition: The maximum resident memory used by all processes of a job (in KB).

Job Run Time

Average Run Time of Done & Exited Jobs
Average Run Time of Done Jobs
Average Run Time of Exited Jobs
Total Run Time of Done & Exited Jobs
Total Run Time of Done Jobs
Total Run Time of Exited Jobs
Definition: The elapsed time from job dispatch to job completion (in seconds)

Job Swap Space Usage

Average Swap Usage of Done & Exited Jobs
Average Swap Usage of Done Jobs
Average Swap Usage of Exited Jobs
Average Swap Usage of Running & Suspended Jobs
Total Swap Usage of Done & Exited Jobs
Total Swap Usage of Done Jobs
Total Swap Usage of Exited Jobs
Total Swap Usage of Running & Suspended Jobs
Definition: The maximum virtual memory used by all processes of a job (in KB).

Job Throughput

Job Throughput
Definition: The number of done and exited jobs divided by the time period to finish these jobs (in jobs/hour).

Job Turnaround Time

Average Turnaround Time of Done & Exited Jobs
Average Turnaround Time of Done Jobs
Average Turnaround Time of Exited Jobs
Definition: The elapsed time from job submission to job completion (in seconds).

Job Wait Time

Average Wait Time of Done & Exited Jobs
Average Wait Time of Done Jobs
Average Wait Time of Exited Jobs
Definition: The elapsed time from job submission to job dispatch (in seconds).

Load Index of Hosts

15-second Run Queue Length
1-minute Run Queue Length
15-minute Run Queue Length
Paging Rate
Disk I/O Rate
Number of Login Users
Interactive Idle Time
Available Memory
Available Swap Space
Available /tmp Space
user-specified external load indices
Definition: The built-in indices provided by LIM and user-defined external dynamic numeric resources.

Memory Utilization of Hosts

Memory Utilization
Definition: The amount of currently available memory divided by the total amount of memory in the host (in percentage).

Number of Hosts

Number of Hosts
Number of OK Hosts
Number of Busy Hosts
Number of Full Hosts
Number of Closed Hosts
Number of Available Hosts

Definition: The total number of LSF Batch Server hosts in the cluster and the number of hosts satisfying certain conditions.

Number of Jobs

Number of Done & Exited Jobs
Number of Done Jobs
Number of Exited Jobs
Number of Running, Pending & Suspended Jobs
Number of Running & Suspended Jobs
Number of Pending Jobs
Definition: The total number of jobs in the cluster satisfying certain conditions.

Swap Space Utilization of Hosts

Swap Space Utilization

Definition: The amount of currently available swap space, divided by the total amount of swap space in the host (in percentage).

Usage of Resource Shared among Hosts

specific shared resource

Definition: The total usage of a resource shared by all or some hosts in the cluster.

C. Statistics

For each statistic, the related category (first field on the Y-axis dialog) and class (second field) are shown.

15-second Run Queue Length (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

1-minute Run Queue Length (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

15-minute Run Queue Length (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Available Memory (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Available Swap Space (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Available /tmp Space (description, page 22)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Average CPU Time of Done Jobs (description, page 19)
Job Related | Job CPU Time
Resource Related | Job CPU Time
Queue Related | Job CPU Time
User Related | Job CPU Time
Project Related | Job CPU Time

C Statistics

Average CPU Time of Done & Exited Jobs (description, page 19)
> *Job Related | Job CPU Time*
> *Resource Related | Job CPU Time*
> *Queue Related | Job CPU Time*
> *User Related | Job CPU Time*
> *Project Related | Job CPU Time*

Average CPU Time of Exited Jobs (description, page 19)
> *Job Related | Job CPU Time*
> *Resource Related | Job CPU Time*
> *Queue Related | Job CPU Time*
> *User Related | Job CPU Time*
> *Project Related | Job CPU Time*

Average CPU Time of Running & Suspended Jobs (description, page 22)
> *Job Related | Job CPU Time*
> *Resource Related | Job CPU Time*
> *Queue Related | Job CPU Time*
> *User Related | Job CPU Time*
> *Project Related | Job CPU Time*

Average Memory Usage of Done Jobs (description, page 20)
> *Job Related | Job Memory Usage*
> *Resource Related | Job Memory Usage*
> *Queue Related | Job Memory Usage*
> *User Related | Job Memory Usage*
> *Project Related | Job Memory Usage*

Average Memory Usage of Done & Exited Jobs (description, page 20)
> *Job Related | Job Memory Usage*
> *Resource Related | Job Memory Usage*
> *Queue Related | Job Memory Usage*
> *User Related | Job Memory Usage*
> *Project Related | Job Memory Usage*

Average Memory Usage of Exited Jobs (description, page 20)
> *Job Related | Job Memory Usage*
> *Resource Related | Job Memory Usage*
> *Queue Related | Job Memory Usage*

User Related | Job Memory Usage
Project Related | Job Memory Usage

Average Memory Usage of Running & Suspended Jobs (description, page 22)
Job Related | Job Memory Usage
Resource Related | Job Memory Usage
Queue Related | Job Memory Usage
User Related | Job Memory Usage
Project Related | Job Memory Usage

Average Run Time of Done Jobs (description, page 19)
Job Related | Job Run Time
Queue Related | Job Run Time
User Related | Job Run Time
Project Related | Job Run Time

Average Run Time of Done & Exited Jobs (description, page 19)
Job Related | Job Run Time
Queue Related | Job Run Time
User Related | Job Run Time
Project Related | Job Run Time

Average Run Time of Exited Jobs (description, page 19)
Job Related | Job Run Time
Queue Related | Job Run Time
User Related | Job Run Time
Project Related | Job Run Time

Average Swap Usage of Done Jobs (description, page 20)
Job Related | Job Swap Space Usage
Resource Related | Job Swap Space Usage
Queue Related | Job Swap Space Usage
User Related | Job Swap Space Usage
Project Related | Job Swap Space Usage

Average Swap Usage of Done & Exited Jobs (description, page 20)
Job Related | Job Swap Space Usage
Resource Related | Job Swap Space Usage
Queue Related | Job Swap Space Usage

User Related | Job Swap Space Usage
Project Related | Job Swap Space Usage

Average Swap Usage of Exited Jobs (description, page 20)
Job Related | Job Swap Space Usage
Resource Related | Job Swap Space Usage
Queue Related | Job Swap Space Usage
User Related | Job Swap Space Usage
Project Related | Job Swap Space Usage

Average Swap Usage of Running & Suspended Jobs (description, page 22)
Job Related | Job Swap Space Usage
Resource Related | Job Swap Space Usage
Queue Related | Job Swap Space Usage
User Related | Job Swap Space Usage
Project Related | Job Swap Space Usage

Average Turnaround Time of Done Jobs (description, page 18)
Job Related | Job Turnaround Time
Host Related | Job Turnaround Time
Queue Related | Job Turnaround Time
User Related | Job Turnaround Time
Project Related | Job Turnaround Time

Average Turnaround Time of Done & Exited Jobs (description, page 18)
Job Related | Job Turnaround Time
Host Related | Job Turnaround Time
Queue Related | Job Turnaround Time
User Related | Job Turnaround Time
Project Related | Job Turnaround Time

Average Turnaround Time of Exited Jobs (description, page 18)
Job Related | Job Turnaround Time
Host Related | Job Turnaround Time
Queue Related | Job Turnaround Time
User Related | Job Turnaround Time
Project Related | Job Turnaround Time

Average Wait Time of Done Jobs (description, page 19)
Job Related | Job Wait Time

Queue Related | Job Wait Time
User Related | Job Wait Time
Project Related | Job Wait Time

Average Wait Time of Done & Exited Jobs (description, page 19)
Job Related | Job Wait Time
Queue Related | Job Wait Time
User Related | Job Wait Time
Project Related | Job Wait Time

Average Wait Time of Exited Jobs (description, page 19)
Job Related | Job Wait Time
Queue Related | Job Wait Time
User Related | Job Wait Time
Project Related | Job Wait Time

Batch Job Slot Utilization (description, page 21)
Resource Related | Batch Job Slot Utilization of Hosts
Host Related | Batch Job Slot Utilization of Hosts

Batch Processor Utilization (description, page 21)
Resource Related | Batch Job Slot Utilization of Hosts
Host Related | Batch Job Slot Utilization of Hosts

CPU Utilization (description, page 21)
Resource Related | CPU Utilization of Hosts
Host Related | CPU Utilization of Hosts

Disk I/O Rate (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Interactive Idle Time (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Job Throughput (description, page 18)
Job Related | Job Throughput
Host Related | Job Throughput
Queue Related | Job Throughput

C   Statistics

Memory Utilization (description, page 21)
Resource Related | Memory Utilization of Hosts
Host Related | Memory Utilization of Hosts

Number of Busy Hosts (description, page 23)
Host Related | Number of Hosts

Number of Closed Hosts (description, page 23)
Host Related | Number of Hosts

Number of Done Jobs (description, page 18)
Job Related | Number of Jobs
Host Related | Number of Jobs
Queue Related | Number of Jobs
User Related | Number of Jobs
Project Related | Number of Jobs

Number of Done and Exited Jobs (description, page 18)
Job Related | Number of Jobs
Host Related | Number of Jobs
Queue Related | Number of Jobs
User Related | Number of Jobs
Project Related | Number of Jobs

Number of Exited Jobs (description, page 18)
Job Related | Number of Jobs
Host Related | Number of Jobs
Queue Related | Number of Jobs
User Related | Number of Jobs
Project Related | Number of Jobs

Number of Full Hosts (description, page 23)
Host Related | Number of Hosts

Number of Hosts in the Cluster (description, page 23)
Host Related | Number of Hosts

Number of Login Users (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Number of OK Hosts (description, page 23)
Host Related | Number of Hosts

Num of Pending Jobs (description, page 22)
Job Related | Number of Jobs
Host Related | Number of Jobs
Queue Related | Number of Jobs
User Related | Number of Jobs
Project Related | Number of Jobs

Num of Running, Pending & Suspended Jobs (description, page 22)
Job Related | Number of Jobs
Host Related | Number of Jobs
Queue Related | Number of Jobs
User Related | Number of Jobs
Project Related | Number of Jobs

Num of Running & Suspended Jobs (description, page 22)
Job Related | Number of Jobs
Host Related | Number of Jobs
Queue Related | Number of Jobs
User Related | Number of Jobs
Project Related | Number of Jobs

Number of Unavailable Hosts (description, page 23)
Host Related | Number of Hosts

Paging Rate (description, page 21)
Resource Related | Load Index of Hosts
Host Related | Load Index of Hosts

Swap Space Utilization (description, page 21)
Resource Related | Swap Space Utilization of Hosts
Host Related | Swap Space Utilization of Hosts

Total CPU Time of Done Jobs (description, page 19)
Job Related | Job CPU Time
Resource Related | Job CPU Time
Queue Related | Job CPU Time

User Related | Job CPU Time
Project Related | Job CPU Time

Total CPU Time of Done & Exited Jobs (description, page 19)
Job Related | Job CPU Time
Resource Related | Job CPU Time
Queue Related | Job CPU Time
User Related | Job CPU Time
Project Related | Job CPU Time

Total CPU Time of Exited Jobs (description, page 19)
Job Related | Job CPU Time
Resource Related | Job CPU Time
Queue Related | Job CPU Time
User Related | Job CPU Time
Project Related | Job CPU Time

Total CPU Time of Running & Suspended Jobs (description, page 22)
Job Related | Job CPU Time
Resource Related | Job CPU Time
Queue Related | Job CPU Time
User Related | Job CPU Time
Project Related | Job CPU Time

Total Memory Usage of Done Jobs (description, page 20)
Job Related | Job Memory Usage
Resource Related | Job Memory Usage
Queue Related | Job Memory Usage
User Related | Job Memory Usage
Project Related | Job Memory Usage

Total Memory Usage of Done & Exited Jobs (description, page 20)
Job Related | Job Memory Usage
Resource Related | Job Memory Usage
Queue Related | Job Memory Usage
User Related | Job Memory Usage
Project Related | Job Memory Usage

Total Memory Usage of Exited Jobs (description, page 20)
Job Related | Job Memory Usage

Queue Related | Job Swap Space Usage
User Related | Job Swap Space Usage
Project Related | Job Swap Space Usage

Total Swap Usage of Exited Jobs (description, page 20)
Job Related | Job Swap Space Usage
Resource Related | Job Swap Space Usage
Queue Related | Job Swap Space Usage
User Related | Job Swap Space Usage
Project Related | Job Swap Space Usage

Total Swap Usage of Running & Suspended Jobs (description, page 22)
Job Related | Job Swap Space Usage
Resource Related | Job Swap Space Usage
Queue Related | Job Swap Space Usage
User Related | Job Swap Space Usage
Project Related | Job Swap Space Usage

Index

Index

Index

W

Windows NT
    database utility commands 56

X

X-axis dialog 26

Y

Y-axis dialog 26
Y-axis statistics 17